

STREAMING A SINGLE MEDIA TRACK TO MULTIPLE CLIENTS

5 Inventors: Geetha Srikantan, Aravind Narasimhan, Seth Proctor, Jan Brittenson,
Matthew Shafer & Jonathan Sergent

BACKGROUND

This invention relates to the field of computer systems. More particularly, in a media streaming system in which a track of media has associated metadata 10 that provides timing, offset and other information, an apparatus and methods are provided for streaming a media track to multiple clients with just one copy of the track metadata.

Media streaming systems are designed to stream media programs and events, which may be pre-recorded or live, to client devices (e.g., computers, 15 media players). A client can play a media stream as it is received (i.e., before the stream is complete), thereby allowing fast, or even real-time, enjoyment of the media program or event.

Streamed media is often composed of multiple tracks. For example, an audiovisual program generally includes at least one audio track and at least one 20 video track. Each track includes media of the appropriate type (e.g., audio, video), and may also include metadata that is used to stream the media correctly. A track's metadata may include information for identifying a media segment or sample (or other unit of media) that should be played for a given time index 25 within the program, for determining where that segment or sample is located in the file, etc. A media track's metadata may thus be used by a media streaming server to stream the media track in the correct sequence, with appropriate timing, and so on.

When streaming a media track or program to multiple clients, existing systems usually store separate copies of each track's metadata for each client. More specifically, from a media file containing the media to be streamed, these systems repeatedly extract and store the metadata each time a new client requests 5 the media. When serving a large number of clients, this could require a significant amount of the server's resources (e.g., memory, processor, disk usage), and may limit the number of clients the server can support.

Also, some existing systems typically allocate just one file descriptor, to be shared among all clients receiving a particular media stream, for accessing the 10 media file containing the media being streamed. This can lead to a great deal of contention among the client streams as each one attempts to seek to (i.e., find) and extract a different media segment or sample.

Yet further, when streaming a multi-track media program to a client, existing media streaming systems attempt to maintain synchronization between 15 the tracks so that the appropriate corresponding track media is streamed or played for each time unit of the program, but may be incapable of regaining synchronization if it is lost. For example, if the media for one track is received (e.g., from a storage device) at a slower rate than the media for another track, an existing media streaming system may simply continue streaming the media, even 20 if the tracks get further and further out of synchronization. Other systems may simply halt a media stream if synchronization is lost, without attempting to correct the situation.

SUMMARY

25 Therefore, in one embodiment of the invention, an apparatus and methods are provided for sharing one copy of a media track's metadata for streaming the media track to multiple clients.

In this embodiment, one file track is generated to store the metadata of a single media track. For a media program comprising multiple tracks, a separate file track is generated for each track. For each client stream of the media, separate file track handles are established. Each file track handle acts as an interface
5 between its client stream – different client streams may be at different time indices within the media – and the single instance of media metadata.

In another embodiment of the invention, a single file descriptor is shared among all file track handles associated with a single media track, for accessing the file storing the track's media.

10

DESCRIPTION OF THE FIGURES

FIG. 1 is a block diagram depicting a server configured to stream media in accordance with an embodiment of the present invention.

15 FIG. 2 is a block diagram depicting the use of a single copy of media track metadata to stream the media to multiple clients, in accordance with an embodiment of the invention.

FIG. 3 depicts one configuration of program objects cooperating to stream a media track to multiple clients with a single copy of the track's metadata, in accordance with an embodiment of the present invention.

20 FIGs. 4A-4B comprise a flowchart demonstrating one method of streaming a media track to multiple clients with a single copy of the track's metadata in accordance with an embodiment of the present invention.

FIG. 5 depicts one configuration of program objects for resynchronizing media within a media stream, in accordance with an embodiment of the present
25 invention.

FIG. 6 is a flowchart demonstrating one method of resynchronizing media within a media stream in accordance with an embodiment of the present invention.

5

DETAILED DESCRIPTION

The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of particular applications of the invention and their requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

15 The program environment in which a present embodiment of the invention is executed illustratively incorporates a general-purpose computer or a special purpose device such as a computer server configured to provide data or media streaming services to computers or communications devices of virtually any configuration (e.g., wired, wireless, portable, desktop). Details of such computers and other devices (e.g., processor, memory, data storage, display) are well known and may be omitted for the sake of clarity. Further, embodiments of the invention are described as they may be implemented in an object-oriented programming environment. Suitable variations of embodiments may be implemented using other programming models or frameworks, as will be appreciated by one skilled 20 in the art.

25 It should also be understood that the techniques of the present invention might be implemented using a variety of technologies. For example, the methods

described herein may be implemented in software executing on a computer system, or implemented in hardware utilizing either a combination of microprocessors or other specially designed application specific integrated circuits, programmable logic devices, or various combinations thereof. In particular, the methods described herein may be implemented by a series of computer-executable instructions residing on a storage medium such as a carrier wave, disk drive, or computer-readable medium. Exemplary forms of carrier waves may take the form of electrical, electromagnetic or optical signals conveying digital data streams along a local network or a publicly accessible network such as the Internet.

In one embodiment of the invention a media or data streaming server is configured to stream media or data to multiple clients. A media program or event that is to be streamed to multiple clients includes one or more tracks (e.g., audio, video), each of which may be stored in one or more media files. When a first request for the media program is received, the track's metadata is retrieved and stored in a file track. The metadata may include information regarding: timing (e.g., to synchronize the track's media with other program tracks), locations of media segments or samples (or other units of media) within a media file, hierarchies of media units (e.g., chunks, samples that comprise a chunk), edit tables, etc.

For each client to which the media track is to be streamed, separate file track handles are established to reference metadata within the file track. Thus, instead of storing separate copies of the track metadata for each client stream, each client's file track handle references (e.g., with pointers) the single file track to retrieve timing location, locate media segments in a media file, etc. Further, each file track handle is allocated a file descriptor into the track's media files, for seeking and retrieving media.

Specific embodiments of the invention are described below as they may be implemented for streaming QuickTime media from a UNIX-based computer system, such as a system executing the Solaris™ operating system by Sun Microsystems, Inc. Such embodiments may be modified for use with other types of media and other computer systems, as may be understood from the following descriptions.

An Illustrative Media Streaming Server

Media streaming allows a user to receive and enjoy media content without having to wait until the full program or presentation is downloaded to his or her client device. For example, the user can enjoy a pre-recorded program, or experience a live event in real-time, without waiting until the full program is received.

Media may be streamed in unicast or multicast mode. In unicast mode the streaming server maintains a dedicated connection to each receiving client device, which grants a user great control over his or her stream. He or she may, for example, be able to pause a stream, rewind or fast-forward through the streamed media, or perform other control functions. However, this can lead to inefficient use of bandwidth for a large number of users. In multicast mode, the media streaming server streams a program to multiple users simultaneously, thereby using less bandwidth. This type of streaming is thus comparable to a traditional broadcast, and users have little control over their individual streams. Live events may normally be streamed in multicast mode, because it is more efficient for serving large numbers of users. And, because it is a live event being enjoyed in real-time, there is little need to manipulate the streamed media.

A media streaming server according to a present embodiment of the invention may operate in a “reflection” mode of operation, in which the server

receives a media stream from another streaming system or server (usually in multicast mode), and forwards the media to one or more users (in unicast or multicast mode).

- Streaming real-time media places constraints upon the issuing server,
- 5 because delivery of each frame or other unit of the media must be performed in a specified order and within a certain period of time. Thus, despite the number of clients it serves, a media streaming server must strive to meet the demands of streaming real-time media so that the quality of service to the users does not drop to an unacceptable level. For example, regardless of the type of program (i.e., live
- 10 or pre-recorded) and mode of streaming (i.e., unicast or multicast), streamed media is generally compressed to decrease the bandwidth that it consumes in transit, thus helping to ensure timely delivery of media to a client.

A media streaming server according to one embodiment of the invention is configured to stream QuickTime media and/or other forms of media, in a unicast or multicast mode, over a proprietary or publicly accessible network such as the Internet. Media streams are formatted according to a set of protocols compatible with the transmission medium. In particular, when streaming QuickTime media the server may be configured for RTSP (Real-Time Streaming Protocol) to facilitate a client's control of a media stream, RTP (Real-Time Transport

15 Protocol) to deliver the stream to the client and/or receive media from another source, RTCP (Real-Time Transport Control Protocol) to receive or exchange information regarding the quality of a stream, SDP (Session Description Protocol) to describe media to the client, etc. Other embodiments may be configured for other media protocols.

20

FIG. 1 depicts media streaming server 102, configured to stream QuickTime media according to an embodiment of the invention. In FIG. 1, media streaming server 102 serves clients 110, 112. The media that is streamed to the

clients may comprise a pre-recorded program retrieved from storage device 104 or a real-time program received from server 130 (e.g., as part of multicast broadcast 130a). Media streaming server 102 may thus stream live events (e.g., concerts, news broadcasts, sporting events), movies, documentaries, training videos, educational programs or classes, etc.

Media streaming may require multiple connections between media streaming server 102 and a client. In the embodiment illustrated in FIG. 1, a first connection is made for RTSP (e.g., connection 110a, connection 112a) to allow a client to control a media stream. In particular, a client uses an RTSP connection 10 to send commands to the media streaming server. The media stream commands that a client may submit to the server in this embodiment include commands such as Options, to receive a list of supported commands; Describe, to prompt the server to describe a media program; Setup, to identify desired tracks it would like to receive (wherein each track may be a different media form, such as video, 15 audio, etc.); Play, to play a media track or program; Pause, to temporarily stop streaming; Teardown, to end a stream; etc. Thus, client 110 may, for example, establish RTSP connection 110a with server 102 and issue the Describe command to receive a description of the content and tracks available for streaming. Client 110 may then submit a Setup request for one or more tracks.

When a client issues a Setup command to the server, the server establishes an RTP connection (e.g., connection 110b, connection 112b) and an RTCP connection (e.g., connection 110c, connection 112c) for the selected track(s). When the Play command is received, the server starts streaming media packets to the client through the RTP connection. And, the server and client may exchange 25 RTCP packets, through the RTCP connection, that describe the quality of the stream. When a Teardown command is issued, the server closes its related stream connections with the issuing client.

The various connections employed by the media streaming server may utilize TCP (Transport Control Protocol) sockets for a compatible communication medium through which the server and a client communicate (e.g., the Internet). In other embodiments of the invention, the sockets may be configured according to a different protocol (e.g., HTTP – HyperText Transport Protocol, FTP – File Transfer Protocol).

As already described, media streaming server 102 of FIG. 1 may stream real-time or live media to clients and may also stream pre-recorded media.

Further, in the reflection mode of operation, the media streaming server may 10 redirect to clients media that it receives from another entity, such as a live event, a video camera, a broadcast from another server (e.g., server 130), etc. In this situation media streaming server 102 acts as a client and receives media packets via an RTP connection established with the entity.

Clients 110, 112 are equipped with suitable media players for playing the 15 media streamed from media streaming server 102. For QuickTime media streaming, the clients may operate a QuickTime player such as that available from Apple Computer, Inc. Client computing devices may operate virtually any operating system for which a suitable media player is available (e.g., Solaris, Mac OS, Windows, Linux). Because client devices may have relatively low- 20 bandwidth communication capability (e.g., 56K modem), media streams may be sent with relatively low bit-rates. Higher bit-rates may, of course, be implemented for more robust clients. Clients may identify media to be streamed to them by submitting a URL (Uniform Resource Locator), file name, program name (e.g., name of a movie, song title), etc.

25 U.S. Patent Application Serial No. 09/xxx,xxx, filed on April 6, 2001 and entitled “Method and Apparatus for Handling Events Received at a Server Socket,” discloses a configuration of media streaming server 102 for streaming

media to multiple clients through a limited number of communication sockets, and is hereby incorporated by reference.

Streaming Media to Multiple Clients With a Single File Track

- 5 In one embodiment of the invention, a media streaming server streams media that is accompanied by metadata to facilitate the streaming process. For example, metadata for a QuickTime media track may indicate which unit or piece of media (e.g., audio sample, movie segment) corresponds to a given time index within the media's program, and where to find it within a media file. In this
- 10 embodiment, the metadata may be assembled (e.g., extracted from one or more media files for a given media track) just once, but may be used to facilitate streaming of that track to multiple clients. This scheme is thus more efficient than existing streaming systems in which separate copies of the metadata are assembled for each individual client stream.
- 15 The media data, once extracted with the help of the metadata, is placed into RTP packets and streamed to the clients. The media data may be prepared for streaming in virtually any format that can be handled by a client. In particular, the media may be placed into packets without altering its coding, compression or other attributes.
- 20 Illustratively, a client may request a media stream comprising multiple tracks, or may individually request specific tracks. Each track has its own metadata and may be sent to a client as a separate stream (i.e., multiple tracks may or may not be combined into a single stream). As described above, in one embodiment of the invention the media is sent in RTP packets, in which case the
- 25 streaming may be controlled by a client via RTSP commands and the quality of a stream may be reported via RTCP. In other embodiments of the invention, other suitable protocols are employed.

When the media streaming server is operating in a reflection mode, in which it receives media (as RTP packets) from a source and streams them to clients, metadata for the media may be received with or in advance of the media.

FIG. 2 depicts the use of a singleton file track and multiple file track

- 5 handles to stream a media track to multiple clients without separately extracting or making separate copies of the track's metadata, according to one embodiment of the invention.

In the illustrated embodiment, media program 200 includes audio track 202 and video track 204. Each track includes its indicated type of media for the
10 program (i.e., audio, video) and metadata that, among other things: identifies which media unit (e.g., audio sample, video frame) corresponds to a given time index in the media program; identifies which chunk (e.g., collection of media units) a given unit belongs to or is included in; lists the media units within a given chunk; identifies where a given chunk or media unit is stored within media
15 program file 200; includes edit tables that allow content tools to insert or delete parts of tracks without copying or compacting the entire track (e.g., they allow a track to be fragmented); etc. Illustratively, audio track 202 and video track 204 are stored as separate files (or sets of files). Alternatively, however, different types of tracks or portions of different types of tracks may be stored in a single
20 file.

A media track's metadata may be stored contiguously within the file(s) containing the track's media, may be distributed throughout the track's media file(s), or may even be stored as a separate track.

In accordance with this embodiment of the invention, each track's
25 metadata is extracted and stored in separate file tracks. Thus, the metadata for audio track 202 is assembled as file track 212, while the metadata for video track

204 is stored in file track 214. The different types of metadata within a track may be stored as separate tables, lists, or other data structures.

For each client that a track of the media program is being streamed to, separate file track handles are established. Thus, for client 230a, file track handles 5 222a, 224a are established for accessing the metadata of file track 212 and 214, respectively. Similarly, for client 230b the metadata stored in file track 212 for audio track 202 is accessed by file track handle 222b, while the metadata of file track 214 for video track 204 is accessed by file track handle 224b. Illustratively, each file track handle includes one or more pointers or other references into the 10 associated file track, with which to access the various types of information.

In addition, each file track handle, or set of file track handles for a given client stream, may be allocated a separate file descriptor for accessing the media file(s) that store a track's media. Thus, in the embodiment of FIG. 2, in which audio track 202 and video track 204 are stored in separate files, each of file track 15 handles 222a, 222b, 224a, 224b is allocated its own file descriptor for locating and extracting media for streaming.

While streaming media to a client (e.g., after the client requests a track or program be played), the file track handles for the client's stream access the file tracks to determine the proper sequence of media (i.e., the media unit(s) to stream 20 for each time index of the media). With this information, the file track handles extract the media from the media file(s) so that it can be placed into RTP packets (or packets configured according to another suitable protocol) and streamed to the client via an RTP connection. When a client terminates its stream, the file track handles are deleted.

25 In one embodiment of the invention, a file track is configured to monitor the number of file track handles maintaining references to it. In this embodiment, when all client streams of the related media track have been terminated, the file

track may remove itself and thereby free additional resources for other processes or media tracks.

FIG. 3 depicts several program objects employed in one embodiment of the invention. Appropriate corresponding program modules configured to operate 5 similarly to the illustrated program objects may be generated for computer systems with non-object oriented methods of programming.

In this embodiment, Track 302 represents a class of objects configured to assemble metadata for a media track and allow multiple file track handles to use the metadata on behalf of different client streams. Multiple types of Track objects 10 may be generated from Track 302, such as LiveTrack 320 for assembling metadata for streaming live or real-time event media, and FileTrack 322 for assembling metadata for stored or pre-recorded media. In one alternative embodiment of the invention, one type of track object (e.g., FileTrack 322) may be configured to facilitate streaming of both pre-recorded and live media. In yet 15 another alternative embodiment, live media may not be in the form of tracks and may therefore be served to clients in some other manner.

QuickTime FileTrack 322a and MPEG FileTrack 322b represent FileTrack objects configured to handle metadata for two specific media forms (i.e., QuickTime and MPEG). Objects generated from Track 302 may inherit attributes 20 and methods useful for handling live media, pre-recorded media, and/or other media. Among the inherited attributes, for example, may be one that keeps a count of track handles that reference a given track.

TrackHandle 304 represents a class of objects configured to facilitate the streaming of track media to a client using a shared collection of track metadata. 25 Illustratively, LiveTrackHandle 340 represents a class of TrackHandle objects configured for streams of live media, while FileTrackHandle 342 represents a class of TrackHandle objects configured for pre-recorded or other stored media.

In one alternative embodiment of the invention, one type of track handle object (e.g., FileTrackHandle 342) may be configured for streaming of both pre-recorded and live media. In yet another alternative embodiment, live media may not take the form of a track and may therefore be served to clients in some other manner

5 than described herein.

From FileTrackHandle 342, different types of FileTrackHandle objects may be instantiated for specific forms of media, such as QuickTime FileTrackHandle 342a for QuickTime media and MPEG FileTrackHandle 342b for MPEG media.

10 In this embodiment of the invention, each TrackHandle object derived from TrackHandle 304 includes suitable methods for moving to a certain time index in a media program or track, reading media data into a buffer (for streaming to a client), etc. Thus, TrackHandle objects maintain state information regarding a client's current play position in a media track, and may include one or more

15 buffers for sending media packets to clients and one or more pointers or references to their corresponding Track objects for accessing metadata.

Once a FileTrack object is established for a given media track, subsequent client requests for that track use the established FileTrack rather than assembling the track's metadata again. Thus, FileTrack objects may be configured to

20 instantiate new FileTrackHandle objects for new client streams.

FIG. 4 demonstrates a method of serving streams of media from one track to multiple clients while maintaining only one copy of the media's metadata, according to one embodiment of the invention. In this method, the same media may be streamed to each client, but with different timing. That is, different client streams may, at any given time, be streaming media from different time indices within the media track.

In state 402 of the illustrated procedure, a request is received to stream some pre-recorded media that is not already being streamed. The request may, for example, be received as an RTSP command to setup and play a particular media track. Other forms of requesting media may be used for other protocols or forms
5 of media.

In state 404, the media streaming server creates a Track object to manage shared access to the requested media's metadata. The illustrated procedure indicates the creation of a FileTrack object, which as described above, may be configured for pre-recorded media programs and tracks.

10 In state 406, the FileTrack object extracts and stores the metadata from the media track. The FileTrack object may have to parse one or more files containing the media track in order to assemble the metadata.

15 In state 408, a FileTrackHandle is generated for the new client stream, to map or interface between the stream-specific client context and the track-specific metadata held by the FileTrack object. The FileTrackHandle object may be provided with one or more references to the assembled metadata (e.g., different references for different tables or different types of metadata). In addition, the FileTrackHandle is allocated a file descriptor or similar resource to facilitate its access to the file(s) containing the media (e.g., to find and extract media to be
20 streamed).

In state 410 the requested media is streamed to the new client. The client may issue commands to control the stream – e.g., to rewind or fast forward to locate a particular part of the media, to pause the streaming, etc.

25 In state 412, while the media is being streamed to one or more clients, the media streaming server is alert for any new requests for the same media. If a new request is received, the illustrated procedure returns to state 408 to set up a new

FileTrackHandle object and a new stream. Otherwise, the procedure continues at state 414.

In state 414, the media streaming server determines if all client streams of the media have been terminated. The server may make this determination, for example, whenever a client stream is terminated. Illustratively, when a client stream is terminated, its associated FileTrackHandle object is deleted. If any streams are still in user or established, or if any new streams are in the process of being set up, the procedure returns to state 410. Otherwise, the procedure continues at state 416.

10 In state 416, after all client streams have ended the FileTrack object is deleted and the metadata is removed from memory.

One skilled in the art will appreciate that the procedure depicted in FIG. 4 is just one method of sharing a single set of media track metadata for multiple client media streams. Other suitable methods may be derived from the foregoing 15 description without exceeding the scope of the invention.

Resynchronizing Media During Streaming

In one embodiment of the invention, a system and methods for resynchronizing a media stream are provided. In this embodiment, a media 20 stream is considered to be out of synchronization if media for one or more tracks in the program or event being streamed are not being streamed fast enough to keep pace with the time index of the stream. In other words, a streamed media program's current time index is the time index for which all tracks should be sending corresponding media. If the media for a particular track lags behind the 25 current time index, or no media is being streamed at all for the track, then the stream may be considered out of synchronization. Loss of synchronization may

occur, for example, if a source of a track's media (e.g., a storage device, a live media feed) cannot supply the media at a fast enough rate.

When synchronization is lost, a media streaming server may attempt to resynchronize the media stream by selecting a new, later, media time index at 5 which to resume streaming. The server then requests or awaits media corresponding to the new time index and, until the new time index is reached, the server may cease streaming or may continue streaming media that is ready to be transmitted (e.g., for a track that is still synchronized).

When the new time index is reached, if the necessary media is received the 10 media streaming server commences streaming the corresponding media. If, however, the media corresponding to the new time index is not available when the new time index is reached, the server may make another attempt to resynchronize at another, later, new time index and may allow more time for retrieval of the corresponding media. Attempts to resynchronize a media stream may be 15 performed a configurable number of times, and the server may terminate the stream if all attempts are unsuccessful.

FIG. 5 depicts several program objects that may cooperate to facilitate resynchronization of a media stream in one embodiment of the invention.

Appropriate corresponding program modules configured to operate similarly to 20 the illustrated program objects may be established for computer systems using non-object oriented methods of programming.

In FIG. 5, TrackStream 502 represents a class of objects configured to process a media track's data for streaming to a client. LiveTrackStream 520 represents a TrackStream class of objects configured for streams of live media, 25 while FileTrackStream 522 represents a TrackStream class of objects configured for pre-recorded or other stored media. A TrackStream object may include interfaces for retrieving specific track media (e.g., for a specified time index).

From FileTrackStream 522, different types of FileTrackStream objects may be instantiated or invoked for specific forms of media, such as QuickTime FileTrackStream 522a for QuickTime media and MPEG FileTrackStream 522b for MPEG media.

5 In one embodiment of the invention, program objects depicted in FIG. 5 may cooperate with objects such as those depicted in and described in conjunction with FIG. 3. In particular, in this embodiment a FileTrackStream object may include or reference a FileTrackHandle object for seeking specific locations in a media file, reading data from the file, etc.

10 Stream 504 is an abstraction of a media stream to a client. Different types of media streams may be represented by, for example, LiveStream 540 for a stream of live media and FileStream 542 for a stream of pre-recorded media. For different media protocols or forms, different types of stream objects may be created to handle the streaming. Thus, in FIG. 5, FileRTPStream 542a is
15 configured to stream media to clients through RTP, while FileHTTPStream 542b is configured to stream media through HTTP.

In the illustrated embodiment of the invention, each Stream object includes methods to start, stop, pause and otherwise control a media stream (e.g., in response to client commands). A Stream object has one or more sockets with
20 which to communicate with a client, and interacts with, or may even include, a TrackStream object to retrieve media data for streaming. Illustratively, and because a media stream may comprise multiple tracks, each Stream object may be associated with, or include, multiple TrackStream objects. For example, if FileRTPStream 542a is created to stream to a client a basic audiovisual media
25 program consisting of one audio track and one video track, FileRTPStream 542a may interact with or include separate FileTrackStream objects for each track. As described below, when a track within a media stream gets out of synchronization,

the Stream object controlling the streaming may initiate a resynchronization procedure.

In one embodiment of the invention, a client connects to a media streaming server to request a pre-recorded media program, and a Stream object of 5 the appropriate type or configuration, such as FileRTPStream 542a is generated to control the streaming of the media. FileRTPStream 542a initiates or invokes one or more TrackStream objects of the appropriate type, such as QuickTime FileTrackStream 522a. The FileTrackStream object(s) then create or invoke Track or TrackHandle objects of the appropriate type or configuration (e.g., as 10 depicted in FIG. 3) to access and retrieve media from the program's track(s). In this embodiment, the Stream object controls the streaming (e.g., starting, stopping, pausing) and resynchronizes the stream if necessary. The TrackStream object(s) control the retrieval of track media by initiating appropriating tasks and/or invoking interfaces provided by the Track or TrackHandle objects (e.g., to 15 locate media data, read data from a media file).

FIG. 6 demonstrates the resynchronization of a media stream according to one embodiment of the invention. The media to be streamed in the illustrated procedure is pre-recorded and is therefore retrieved from one or more storage devices, but the procedure may be readily modified for live media received from a 20 live event or other server. In state 602, a pre-recorded media program comprising multiple tracks is being streamed to a client by an assigned Stream object (e.g., FileRTPStream 542a of FIG. 5).

In state 604, a request is made or scheduled (e.g., by QuickTime FileTrackStream 522a of FIG. 5) to read media data for one of the program's 25 tracks from a media file. Illustratively, the requested media data corresponds to an upcoming media time index (e.g., movie time). While waiting for the requested data, the Stream object may stream media data that is already available

(e.g., because it was retrieved earlier), for the same or a different track, or may wait until the requested data should be available.

In state 606, the Stream object determines whether the requested data is available, or may be automatically woken or otherwise alerted when data is received. If the requested data has been received, then the procedure returns to state 602 to stream the media data at the appropriate media time index. In particular, if media data for the program's tracks is received in a timely fashion, the Stream object may stream the data at the corresponding media time index, thereby keeping the media program synchronized. If the Stream object determines in state 606 that the requested data is not available (e.g., a storage device is over-burdened), thus indicating a loss of synchronization, the procedure advances to state 608.

In state 608, the Stream object advances or increments the current media time index to a future time index and, in state 610, requests media data corresponding to that future index. Thus, in this embodiment of the invention, the media streaming server opts to drop or ignore data corresponding to the media time between the current and future time indices. Illustratively, any such data (e.g., for a track that has not lost synchronization) may be discarded. The Stream object then goes to sleep or waits an appropriate amount of time for the requested media to be received.

In state 612, the Stream object determines whether the media data corresponding to the new media time index has been received and is available for streaming. If so, the Stream object returns to state 602 to stream the media and retrieve the next portion of media.

However, if the media data is not received and available for streaming at the new time index, in state 614 the Stream object determines whether it has failed to resynchronize a threshold number of times (e.g., three). If not, the illustrated

procedure returns to state 608 to advance the media time index once again. Each successive time that the time index is advanced, it may be advanced by a larger increment. For example, if the time index is advanced by a period T the first time, the next time it may be advanced $2*T$, then $4*T$, $8*T$, etc. If resynchronization 5 fails a threshold number of times, the Stream object halts the stream and the procedure ends.

One skilled in the art will appreciate that the procedure depicted in FIG. 6 is just one method of resynchronizing a media stream by dropping old media in favor of regaining synchronization at a later time index. Other suitable methods 10 may be derived from the foregoing description without exceeding the scope of the invention.

The amount of media data requested or retrieved each time it is needed may be determined in part by the size or number of available data buffers, the frequency with which data requests are issued, the bit rate at which the media is 15 streamed to a client, etc.

The foregoing descriptions of embodiments of the invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the invention to the forms disclosed. Many modifications and variations will be apparent to practitioners skilled in the art. 20 Accordingly, the above disclosure is not intended to limit the invention; the scope of the invention is defined by the appended claims.